arXiv:1512.07934v1 [math.ST] 25 Dec 2015

# A SCALABLE QUASI-BAYESIAN FRAMEWORK FOR GAUSSIAN GRAPHICAL MODELS

YVES F. ATCHADÉ

(Dec. 2015)

Abstract. This paper deals with the Bayesian estimation of high dimensional Gaussian graphical models. We develop a quasi-Bayesian implementation of the neighborhood selection method of Meinshausen and Buhlmann (2006) for the estimation of Gaussian graphical models. The method produces a product-form quasi-posterior distribution that can be efficiently explored by parallel computing. We derive a non-asymptotic bound on the contraction rate of the quasi-posterior distribution. The result shows that the proposed quasi-posterior distribution contracts towards the true precision matrix at a rate given by the worst contraction rate of the linear regressions that are involved in the neighborhood selection. We develop a Markov Chain Monte Carlo algorithm for approximate computations, following an approach from Atchadé (2015a). We illustrate the methodology with a simulation study. The results show that the proposed method can fit Gaussian graphical models at a scale unmatched by other Bayesian methods for graphical models.

## 1. Introduction

We consider the problem of fitting large Gaussian graphical models with diverging number of parameters from limited data. This amount to estimating a sparse precision matrix $\vartheta \in \mathcal{M}_p^+$ from $p$-dimensional Gaussian observations $y^{(i)} \in \mathbb{R}^p$, $i = 1, \ldots, n$, where $\mathcal{M}_p^+$ denotes the cone of $\mathbb{R}^{p\times p}$ of symmetric positive definite matrices. The frequentist approach to this problem has generated an impressive literature over the last decade or so (see for instance Bühlmann and van de Geer (2011); Hastie et al. (2015) and the reference therein).

There is currently an interest, particularly in biomedical research, for statistical methodologies that can allow practitioners to incorporate external information in fitting such graphical models (Mukherjee and Speed (2008); Peterson et al. (2015)). This

2000 *Mathematics Subject Classification.* 60F15, 60G42.

*Key words and phrases.* Gaussian graphical models, Quasi-Bayesian inference, Pseudo-likelihood, Posterior contraction, Moreau-Yosida approximation, Markov Chain Monte Carlo.

Y. F. Atchadé: University of Michigan, 1085 South University, Ann Arbor, 48109, MI, United States. *E-mail address:* yvesa@umich.edu.

problem naturally calls for a Bayesian formulation (Dobra et al. (2011); Lenkoski and Dobra (2011); Wang and Li (2012); Khondker et al. (2013); Peterson et al. (2015)). However, most existing Bayesian methods for fitting graphical models do not scale well with $p$, the number of nodes in the graph. The main difficulty is computational, and hinges on the ability to handle interesting prior distributions on $\mathcal{M}_p^+$. The most commonly used class of priors distributions for Gaussian graphical models is the class of G-Wishart distributions (Atay-Kayis and Massam (2005)). However G-Wishart distributions have intractable normalizing constants, and become impractical for inferring large graphical models, due to the cost of approximating the normalizing constants (Dobra et al. (2011); Lenkoski and Dobra (2011); Wang and Li (2012)). Following the development of the Bayesian lasso of Park and Casella (2008) and other Bayesian shrinkage priors for linear regressions (Carvalho et al. (2010)), several authors have proposed prior distributions on $\mathcal{M}_p^+$ obtained by putting conditionally independent shrinkage priors on the entries of the matrix, subject to a positive definiteness constraint (Wang (2012); Khondker et al. (2013)). The main drawback of this approach is that these prior distributions are constructed explicitly so as to cancel the intractable normalizing constants, raising the issue of the impact of such prior distribution trick on the inference. Furthermore, dealing with the positive definiteness constraint in the posterior distribution requires careful MCMC design, and becomes a limiting factor for large $p$. So it appears that most existing Bayesian methods for high-dimensional graphical models do not scale well, and can fit only small to moderately large models (upto $p = 200$).

Building on some recent works Atchadé (2015a,b), we develop a quasi-Bayesian approach for fitting large Gaussian graphical models. Our general approach to the problem consists in working with a "larger" pseudo-model $\{\check{f}_\theta,\ \theta \in \check{\Theta}\}$, where $\mathcal{M}_p^+ \subseteq \check{\Theta}$. By pseudo-model we mean that the function $z \mapsto \check{f}_\theta(z)$ is typically not a density on $\mathbb{R}^p$, but $\check{f}_\theta$ is chosen such that the function $\theta \mapsto \sum_{i=1}^n \log \check{f}_\theta(y^{(i)})$ is a good candidate for M-estimation of $\vartheta$. The enlargement of the model space from $\mathcal{M}_p^+$ to $\check{\Theta}$ allows us to relax the positive definiteness constraint. With a prior distribution $\Pi$ on $\check{\Theta}$, we obtain a quasi-posterior distribution $\check{\Pi}_{n,p}$ (not a proper posterior distribution), since $\check{f}_\theta$ is not a proper likelihood function. In the specific case of Gaussian graphical models, we propose to take $\check{\Theta}$ as the space of matrices with positive diagonals, and to take $z \mapsto \check{f}_\theta(z)$ as the pseudo-model underpinning the neighborhood selection method of Meinshausen and Buhlmann (2006). This choice gives a quasi-posterior distribution $\check{\Pi}_{n,p}$ that factorizes, and leads to a drastic improvement in the computing time needed for MCMC computation when a parallel computing architecture is used. We illustrate

the method in Section 4 using simulated data where the number of nodes in the graph is $p \in \{100, 500, 1000\}$.

The idea of replacing the likelihood function by a pseudo-likelihood (or quasi-likelihood) function in a Bayesian inference is not new and has been developed in other contexts, such as in Bayesian semi-parametric inference (Kato (2013); Li and Jiang (2014), and the references therein), and in approximate Bayesian computation (Fearnhead and Prangle (2010)). A general analysis of the contraction properties of these distributions for high-dimensional problems can be found in Atchadé (2015b).

We study the contraction properties of the quasi-posterior distribution $\check{\Pi}_{n,p}$ as $n, p \to \infty$. Under the assumption that there exists a true precision matrix, and some additional assumption on the prior distribution, we show that when the true precision matrix is well conditioned, then $\check{\Pi}_{n,p}$ contracts[1] at the rate $\sqrt{\frac{\bar{s}\log(p)}{n}}$ (see Theorem 7 for a precise statement), where $\bar{s}$ can be viewed as an upper-bound on the largest degree in the un-directed graph defined by the true precision matrix. This convergence rate corresponds to the worst convergence rate that we get from the Bayesian analysis of the linear regressions involved in the neighborhood selection. The condition on the sample size $n$ for the results mentioned above to hold is $n = O(\bar{s}\log(p))$, which shows that the quasi-posterior distribution can concentrate around the true value, even in cases where $p$ exceeds $n$.

The rest of the paper is organized as follows. Section 2 provides a general discussion of quasi-models and quasi-Bayesian inference. The section ends with the introduction of the proposed quasi-Bayesian distribution, based on the neighborhood selection method of Meinshausen and Buhlmann (2006). We specialized the discussion to Gaussian graphical models in Section 3. The theoretical analysis focuses on the Gaussian case, and is presented in Section 3, but the proofs are postponed to Section 6. The simulation study is presented in Section 4. We end the paper with some concluding thoughts in Section 5. A `MATLAB` implementation of the method is available from the author's website.

## 2. Quasi-Bayesian inference of graphical models

For integer $p \geq 1$, and $i = 1, \ldots, p$, let $\mathsf{Y}_i$ be a nonempty subset of $\mathbb{R}$, and set $\mathsf{Y} \stackrel{\text{def}}{=} \mathsf{Y}_1 \times \cdots \times \mathsf{Y}_p$, that we assume is equipped with a reference sigma-finite product measure $\mathrm{d}y$. We first consider a class of Markov random field distributions $\{f_\omega,\ \omega \in \Omega\}$ for joint modeling of $\mathsf{Y}$-valued random variables. We discuss several quasi-Bayesian methods for fitting such models, including the proposed method

[1]The contraction rate is measured in the $L_{\infty,2}$ matrix norm, defined as the largest $L_2$ column norms

based on the neighborhood selection method of Meinshausen and Buhlmann (2006). Gaussian graphical models are then discussed in more detail as special case in Section 3.

let $\mathcal{M}_p$ denote the set of all real symmetric $p \times p$ matrices equipped with the inner product $\langle A, B\rangle_{\mathsf{F}} \stackrel{\text{def}}{=} \sum_{i\leq j} A_i B_{ij}$, and norm $\|A\|_{\mathsf{F}} \stackrel{\text{def}}{=} \sqrt{\langle A, A\rangle_{\mathsf{F}}}$. As above, $\mathcal{M}_p^+$ denotes the subset of $\mathcal{M}_p$ of positive definite matrices. For $i = 1, \ldots, p$, and $1 \leq j < k \leq p$, let $B_i : \mathsf{Y}_i \to \mathbb{R}$ and $B_{jk} : \mathsf{Y}_j \times \mathsf{Y}_k \to \mathbb{R}$ be non-zero measurable functions that we assume known. From these functions we define a $\mathcal{M}_p$-valued function $\bar{B} : \mathsf{Y} \to \mathcal{M}_p$ by

$$(\bar{B}(y))_{ij} = \begin{cases} B_i(y_i) & \text{if } i = j, \\ B_{ij}(y_i, y_j) & \text{if } i < j, \\ B_{ji}(y_j, y_i), & \text{if } j < i. \end{cases}$$

These functions define the parameter space

$$\Omega \stackrel{\text{def}}{=} \left\{ \omega \in \mathcal{M}_p : \; Z(\omega) \stackrel{\text{def}}{=} \int_{\mathsf{Y}} e^{-\langle \omega, \bar{B}(y)\rangle_{\mathsf{F}}} \mathrm{d}y < \infty \right\}.$$

We assume that $\mathsf{Y}$ and $\bar{B}$ are such that $\Omega$ is non-empty, and we consider the exponential family $\{f_\omega, \; \omega \in \Omega\}$ of densities $f_\omega$ on $\mathsf{Y}$ given by

$$f_\omega(y) = \exp\left( - \left\langle \omega, \bar{B}(y) \right\rangle_{\mathsf{F}} - \log Z(\omega) \right), \quad y \in \mathsf{Y}. \tag{1}$$

The model $\{f_\omega, \; \omega \in \Omega\}$ can be useful to capture the dependence structure between a set of $p$ random variables taking values in $\mathsf{Y}$. The version posited in (1) can accommodate a mix of discrete and continuous measurements. The functions $B_i$ are typically viewed as describing the marginal behaviors of the observations in the absence of dependence. Whereas the functions $B_{ij}$ govern the interactions. These marginal and interaction functions are modulated by the parameter $\omega$. More precisely, if $(Y_1, \ldots, Y_p) \sim f_\omega$, then the parameter $\omega$ encodes the conditional independence structure among the $p$ variables $(Y_1, \ldots, Y_p)$. In particular for $i \neq j$, $\omega_{ij} = 0$ means that $Y_i$ and $Y_j$ are conditionally independent given all other variables. The random variables $(Y_1, \ldots, Y_p)$ can then be represented by an undirect graph where there is an edge between $i$ and $j$ if and only if $\omega_{ij} \neq 0$. This type of models are very useful in practice to tease out direct and indirect connections between sets of random variables.

**Example 1** (Gaussian graphical models)**.** One recovers the Gaussian graphical model by taking $\mathsf{Y}_i = \mathbb{R}$, $B_i(x) = x^2/2$, $B_{ij}(x, y) = xy$, $i < j$. In this case $\mathsf{Y} = \mathbb{R}^p$ equipped with the Lebesgue measure, and $\Omega = \mathcal{M}_p^+$.

**Example 2** (Potts models)**.** For integer $M \geq 2$, one recovers the $M$-states Potts model by taking $\mathsf{Y}_i = \{1, \ldots, M\}$, $B_i(x) = x$, and $B_{ij}(x, y) = \mathbf{1}_{\{x=y\}}$. In this case,

$\mathsf{Y} = \{1, \ldots, M\}^d$ equipped with the counting measure. Since $\mathsf{Y}$ is a finite set, we have $\Omega = \mathcal{M}_p$. An important special case of the Potts model is a version of the Ising model where $M = 2$.

Beyond these examples commonly used in the statistics literature, the family $\{f_\omega,\ \omega \in \Omega\}$ also include the class of auto-models proposed by J. Besag (Besag (1974)), the mixed discrete-continuous graphical models proposed in (Cheng et al. (2013); Yang et al. (2014)), as well as few other models used in machine learning, such as Boltzmann machines and Hopefield models (Salakhutdinov and Hinton (2009)).

Suppose that we observe data $y^{(1)}, \cdots, y^{(n)}$ where $y^{(i)} = (y_1^{(i)}, \ldots, y_p^{(i)})' \in \mathsf{Y}$ is viewed as a column vector. We set $x \overset{\text{def}}{=} [y^{(1)}, \ldots, y^{(n)}]' \in \mathbb{R}^{n \times p}$. Given a prior distribution $\Pi$ on $\Omega$, and given the data $x$, the resulting posterior distribution for learning $\omega$ is

$$\Pi_n(A|x) = \frac{\int_A \prod_{i=1}^n f_\omega(y^{(i)}) \Pi(\mathrm{d}\omega)}{\int_\Omega \prod_{i=1}^n f_\omega(y^{(i)}) \Pi(\mathrm{d}\omega)}, \quad A \subseteq \Omega.$$

However, and as discussed in the introduction, posterior distributions from Markov random fields are typically doubly-intractable[2]. There has been some recent advances in MCMC methodology to deal with doubly-intractable distributions (see Lyne et al. (2013) and the references therein). However most of these MCMC algorithms do not scale well to high-dimensional parameter spaces.

In the frequentist literature, a commonly used approach to circumvent computational difficulties with graphical models consists in replacing the likelihood function by a pseudo-likelihood function. For $\omega \in \mathcal{M}_p$, let $\omega_{\cdot i}$ denote the $i$-th column of $\omega$. Note that in the present case, if $(Y_1, \ldots, Y_p) \sim f_\omega$, then for $1 \leq j \leq p$, the conditional distribution of $Y_j$ given $\{Y_k,\ k \neq j\}$ depends on $\omega$ only through the $j$-th column $\omega_{\cdot j}$. We write this conditional distribution as $u \mapsto f^{(j)}_{\omega_{\cdot j}}(u|y_{-j})$, where for $y \in \mathsf{Y}$, $y_{-j} \overset{\text{def}}{=} (y_1, \ldots, y_{j-1}, y_{j+1}, \ldots, y_p)$, (with obvious modifications when $j = 1, p$). Let

$$\tilde{\Omega} \overset{\text{def}}{=} \Big\{\omega \in \mathcal{M}_p :\ u \mapsto f^{(j)}_{\omega_{\cdot j}}(u|y_{-j}) \text{ is a well-defined density on } \mathsf{Y}_j, \\ \text{for all } y \in \mathsf{Y}, \text{ and all } 1 \leq j \leq p\Big\}.$$

Note that $\Omega \subseteq \tilde{\Omega}$. The most commonly used pseudo-likelihood method consists in replacing the initial likelihood contribution $f_\omega(y^{(i)})$ by

$$\tilde{f}_\omega(y^{(i)}) = \prod_{j=1}^p f^{(j)}_{\omega_{\cdot j}}(y_j^{(i)}|y_{-j}^{(i)}), \quad \omega \in \tilde{\Omega}. \tag{2}$$

[2]a terminology introduced by Murray et al. (2006) to mean that the expression of the distribution depends on terms (typically normalizing constants) that cannot be explicitly computed

This pseudo-likelihood approach, which can be viewed as replacing the model $\{f_\omega,\ \omega \in \Omega\}$ by the pseudo-model $\{\tilde{f}_\omega,\ \omega \in \tilde{\Omega}\}$, typically brings important simplifications. For instance, in the Gaussian case, the parameter space $\tilde{\Omega}$ corresponds to the space of symmetric matrices with positive diagonals elements, which has a simpler geometry compared to $\mathcal{M}_p^+$. And in the case of discrete graphical models, the conditional models typically have tractable normalizing constants. Despite the fact that $\{\tilde{f}_\omega,\ \omega \in \tilde{\Omega}\}$ is not a proper statistical model, the quasi-likelihood function $\omega \mapsto \sum_{i=1}^n \log \tilde{f}_\omega(y^{(i)})$ still typically leads to consistent estimates of the parameter. The idea goes back to Besag (1974), and penalized versions of pseudo-likelihood functions have been employed by several authors to fit high-dimensional graphical models.

A closely related idea is the generalized method of moments (GMM). Given $\omega \in \Omega$, $1 \leq j \leq p$, and $y \in \mathsf{Y}$, define

$$m^{(j)}(\omega_{\cdot j}; y_{-j}) \stackrel{\text{def}}{=} \int_{\mathsf{Y}_j} u f^{(j)}_{\omega_{\cdot j}}(u|y_{-j}) \mathrm{d}u.$$

Suppose for instance that these conditional moments are well-defined and available in closed form. Then one can derive another pseudo-model $\{\bar{f}_\omega,\ \omega \in \bar{\Omega}\}$, by taking

$$\bar{\Omega} \stackrel{\text{def}}{=} \left\{\omega \in \mathcal{M}_p :\ \int_{\mathsf{Y}_j} |u| f^{(j)}_{\omega_{\cdot j}}(u|y_{-j}) \mathrm{d}u < \infty \text{ for all } y \in \mathsf{Y}, \text{ and all } 1 \leq j \leq p\right\}.$$

and

$$\bar{f}_\omega(y^{(i)}) \stackrel{\text{def}}{=} \exp\left(-\sum_{j=1}^p \frac{1}{2\sigma_j^2}\left(y_j^{(i)} - m^{(j)}(\omega_{\cdot j}; y_{-j}^{(i)})\right)^2\right), \tag{3}$$

for positive constants $\sigma_j^2$, $j = 1, \ldots, p$. We note that if all the conditional moments of densities in $\{f_\omega,\ \omega \in \Omega\}$ are well defined, then $\Omega \subseteq \bar{\Omega}$. The function (3) is the GMM objective function associated with the moment restrictions

$$\mathbb{E}_\omega\left[Y_j - m^{(j)}(\omega_{\cdot j}; Y_{-j})\right] = 0,\ j = 1, \ldots, p.$$

In the Gaussian case the two pseudo-likelihood functions (2) and (3) coincide. The moment restriction approach is however more flexible in terms of distributional assumptions.

Another method for deriving a pseudo-model for this problem is suggested by the neighborhood selection of Meinshausen and Buhlmann (2006). The idea consists in relaxing the symmetry constraint in $\tilde{\Omega}$. For $1 \leq j \leq p$, we set

$$\Omega_j \stackrel{\text{def}}{=} \Big\{\theta \in \mathbb{R}^p :\ u \mapsto f^{(j)}_\theta(u|y_{-j}) \text{ is a well-defined density on } \mathsf{Y}_j, \text{ for all } y \in \mathsf{Y}, \text{ and all } 1 \leq j \leq p\Big\}.$$

We note that if $\omega \in \Omega$, then $\omega_{\cdot j} \in \Omega_j$. Hence these sets $\Omega_j$ are nonempty, and we define $\check{\Omega} \stackrel{\text{def}}{=} \Omega_1 \times \cdots \times \Omega_p$, that we identify as a subset of the space of $p \times p$ real matrices $\mathbb{R}^{p \times p}$. In particular if $\omega \in \check{\Omega}$, and consistently with our notation above, $\omega_{\cdot j}$ denotes the $j$-column of $\omega$. We consider the pseudo-model $\{\check{f}_\omega,\ \omega \in \check{\Omega}\}$, where

$$\check{f}_\omega(y) \stackrel{\text{def}}{=} \prod_{j=1}^{p} f^{(j)}_{\omega_{\cdot j}}(y_j|y_{-j}), \quad \omega \in \check{\Omega}, \quad y \in \mathsf{Y}. \tag{4}$$

Notice that by definition $\check{\Omega}$ is a product space, whereas $\tilde{\Omega}$ is not, due to the symmetry constraint. This implies that $\omega \mapsto \check{f}_\omega(y)$ factorizes along the columns of $\omega$, whereas $\omega \mapsto \tilde{f}_\omega(y)$ typically does not. One can then maximize a penalized version of $\omega \mapsto \sum_{i=1}^n \log \check{f}_\omega(y^{(i)})$, and this corresponds to the neighborhood selection method of (Meinshausen and Buhlmann (2006), see also Sun and Zhang (2013)). The optimization can be advantageously solved in parallel for each component if the penalty is separable.

As it turns out, all these pseudo-models can also be used in the Bayesian framework, as shown by the seminal work of Chernozhukov and Hong (2003). We shall focus on the pseudo-model (4). With a prior distribution $\Pi$ on $\check{\Omega}$, the quasi-likelihood function $\omega \mapsto \check{f}_\omega$ leads to a quasi-posterior distribution given by

$$\check{\Pi}_{n,p}(A|x) = \frac{\int_A \prod_{i=1}^n \check{f}_\omega(y^{(i)})\Pi(\mathrm{d}\omega)}{\int_{\check{\Omega}} \prod_{i=1}^n \check{f}_\omega(^{(i)})\Pi(\mathrm{d}\omega)}, \quad A \subset \check{\Omega},$$

for which MCMC algorithms can be constructed. Let us assume that the prior distribution factorizes: $\Pi(\mathrm{d}\omega) = \prod_{j=1}^p \Pi_j(\omega_{\cdot j})$. Then we are led to the quasi-posterior distribution

$$\check{\Pi}_{n,p}(\mathrm{d}u_1, \cdots \mathrm{d}u_p|x) = \prod_{j=1}^{p} \check{\Pi}_{n,p,j}(\mathrm{d}u_j|x), \tag{5}$$

where

$$\check{\Pi}_{n,p,j}(\mathrm{d}u|x) = \frac{\prod_{i=1}^n f^{(j)}_u(y^{(i)}_j|y^{(i)}_{-j})\Pi_j(\mathrm{d}u)}{\int_{\Omega_j} \prod_{i=1}^n f^{(j)}_u(y^{(i)}_j|y^{(i)}_{-j})\Pi_j(\mathrm{d}u)},$$

is a probability measure on $\Omega_j$. Basically, relaxing the symmetry allows us to factorize the quasi-likelihood function and this leads to a factorized quasi-posterior distribution, as in (5). Each component of this quasi-posterior distribution can then be explored independently. Despite its simplicity, when used in a parallel computing environment, this approach increases by one order of magnitude the size of graphical models that can be estimated.

**Remark 3** (symmetrization and positive definiteness)**.** One of the limitation of the quasi-Bayesian approach outlined above is that the distribution $\check{\Pi}_{n,p}$ does not necessarily produce symmetric and positive definite matrices. However, because of the contraction properties of $\check{\Pi}_{n,p}$ discussed below, when the true precision matrix is well conditioned, typical realizations of $\check{\Pi}_{n,p}$ are actually symmetric and positive definite, with high probability. From a practical viewpoint, one can remedy a broken symmetry using various symmetrization rules as suggested for instance in Meinshausen and Buhlmann (2006). Lack of positive definiteness is more expensive to repair, but can be addressed for instance by projection of the convex cone of semipositive definite matrices via eigendecomposition (Higham (1988)), and by addition of a small diagonal matrix.

## 3. Gaussian graphical models

We now specialize the discussion to the Gaussian case, where $\mathsf{Y}_i = \mathbb{R}$, $B_i(x) = x^2/2$, and $B_{ij}(x,y) = xy$. Hence in this case, $\Omega = \mathcal{M}_p^+$, $\tilde{\Omega}$ corresponds to the set of symmetric matrices with positive diagonal elements, which is an important simplification over $\mathcal{M}_p^+$. Further dropping the symmetry leads to $\check{\Omega}$, which here is the space of $p \times p$ real matrices with positive diagonal. Assuming that the diagonal elements are known and given, we shall identify $\check{\Omega}$ with the matrix space $\mathbb{R}^{(p-1)\times p}$.

If $\vartheta \in \mathcal{M}_p^+$, and $(Y_1, \ldots, Y_p) \sim f_\vartheta$, it is well known that for all $j \in \{1, \ldots, p\}$, the conditional distribution of $Y_j$ given $Y_k = y_k$, for $k \neq j$ is

$$\mathbf{N}\left( -\sum_{k \neq j} \frac{\vartheta_{kj}}{\vartheta_{jj}} y_k, \frac{1}{\vartheta_{jj}} \right), \tag{6}$$

where $\mathbf{N}(\mu, \sigma^2)$ denotes the Gaussian distribution with mean $\mu$ and variance $\sigma^2$. Given data $x \in \mathbb{R}^{n\times p}$, given $\sigma_j^2 > 0$, and given these conditional distributions, the product of the quasi-model (4) across the data set gives (upto normalizing constants that we ignore) the quasi-likelihood

$$q(\theta; x) \stackrel{\text{def}}{=} \prod_{j=1}^{p} q_j(\theta_{\cdot j}; x),$$

$$\text{with } \; q_j(\theta_{\cdot j}; x) \stackrel{\text{def}}{=} \exp\left( -\frac{1}{2\sigma_j^2} \|x_{\cdot j} - x^{(j)} \theta_{\cdot j}\|_2^2 \right), \;\; \theta \in \mathbb{R}^{(p-1)\times p}, \tag{7}$$

where $x^{(j)} \in \mathbb{R}^{n\times(p-1)}$ is the matrix obtained from $x$ by removing the $j$-th column, and $x_{\cdot j}$ (resp. $\theta_{\cdot j}$) denotes the $j$-column of $x$ (resp. $\theta$). Given (6), it is clear that $\sigma_j^2$ is a proxy for $1/\vartheta_{jj}$. It is also clear that maximizing (7) or a penalized version thereof would give an estimate of $-(\vartheta_{kj}/\vartheta_{jj})_{k\neq j}$. This is precisely the idea of the

neighborhood selection of Meinshausen and Buhlmann (2006), or the sparse matrix inversion with scaled lasso of Sun and Zhang (2013). These methods can be used to recover the sign (the structure) of $\vartheta$, but also gives an estimate of $\vartheta_{kj}$ if $\sigma_j^2$ is a good estimate of $1/\vartheta_{jj}$, or if $\vartheta_{jj}$ can also be estimated (as in the case of the scaled-lasso). We combine (7) with a prior distribution $\Pi(\mathrm{d}\theta) = \prod_{j=1}^p \Pi_j(\mathrm{d}\theta_{\cdot j})$ to obtain a quasi-posterior distribution on $\mathbb{R}^{(p-1)\times p}$ given by

$$\check{\Pi}_{n,p}(\mathrm{d}\theta|x) = \prod_{j=1}^{p} \check{\Pi}_{n,p,j}(\mathrm{d}\theta_{\cdot j}|x, \sigma_j^2), \tag{8}$$

where $\check{\Pi}_{n,p,j}(\cdot|x,\sigma_j^2)$ is the probability measure on $\mathbb{R}^{p-1}$ given by

$$\check{\Pi}_{n,p,j}(\mathrm{d}z|x,\sigma_j^2) \propto q_j(z;x)\Pi_j(\mathrm{d}z).$$

Again the main appeal of $\check{\Pi}_{n,p}$ is its factorized form, which implies that Monte Carlo samples from $\check{\Pi}_{n,p}$ can be obtained by sampling in parallel from the $p$ distributions $\check{\Pi}_{n,p,j}$.

3.1. **Prior distribution.** We address here the choice of the prior distribution $\Pi_j$. Since we are dealing with a linear regression problem, there are many possible ways to set up the prior. We advocate the use of discrete-continuous mixture distributions because these prior distributions have well-understood posterior contraction properties (Castillo et al. (2014); Atchadé (2015b)), and are known to produce sparse posterior samples.

For each $j \in \{1,\ldots,p\}$, we build the prior $\Pi_j$ on $\mathbb{R}^{(p-1)}$ as in Atchadé (2015b). First, let $\Delta_p \stackrel{\text{def}}{=} \{0,1\}^{p-1}$, and let $\{\pi_\delta,\ \delta \in \Delta_p\}$ denote a discrete probability distribution on $\Delta_p$ (which we assume to be the same for all the components $j$). We take $\Pi_j$ as the distribution of the random variable $u \in \mathbb{R}^{p-1}$ obtained as follows.

$$\delta \sim \{\pi_\delta,\ \delta \in \Delta_p\}.\ \text{Given}\ \ \delta,\ (u_1,\ldots,u_{p-1})\ \ \text{are conditionally independent}$$
$$\text{and}\ \ u_k|\delta \sim \begin{cases} \mathsf{Dirac}(0) & \text{if } \delta_k = 0 \\ \mathsf{EN}(\rho_{1j},\rho_{2j}) & \text{if } \delta_k = 1 \end{cases}, \tag{9}$$

where $\mathsf{Dirac}(0)$ is the Dirac measure on $\mathbb{R}$ with mass at 0, and $\mathsf{EN}(\rho_{1j},\rho_{2j})$ denotes the elastic net distribution on $\mathbb{R}$ with density proportional to

$$z \mapsto \frac{1}{C_j}\exp\left(-\alpha\frac{\rho_{1j}}{\sigma_j^2}|z| - (1-\alpha)\frac{\rho_{2j}}{\sigma_j^2}\frac{z^2}{2}\right),\ \ z \in \mathbb{R}, \tag{10}$$

for parameters $\rho_{1j}, \rho_{2j} > 0$, and where $\alpha \in (0, 1]$ is a fixed parameter (in the simulation we use $\alpha = 0.9$). The term $C_j$ is the normalizing constant[3]. We use a fully-Bayesian approach for selecting $\rho_{1j}, \rho_{2j}$. More precisely, we assume that $\rho_{1j}$ and $\rho_{2j}$ have independent prior distribution $\rho_{1j} \sim \phi$, $\rho_{2j} \sim \phi$, where $\phi$ is the uniform distribution $\mathbf{U}(a_1, a_2)$ for $a_1 = 10^{-5}$, and where the choice of $a_2$ follows Atchadé (2015a) Section 4.

We focus on situations where, although $p$ is possibly large, the undirected graph defined by the underlying precision matrix is sparse. This prior information is encoded in the prior distribution, by choosing $\pi_\delta$ as follows. We assume that the components of $\delta$ are conditionally independent with distribution $\mathbf{Ber}(\mathsf{q})$ given $\mathsf{q}$, where $\mathsf{q} \sim \mathbf{Beta}(1, p^u)$, for some $u > 1$. Hence according to the prior distribution, the proportion of non-zero component of each column of $\theta$ is $1/p^{u-1}$. We use $u = 1.5$.

With the prior distribution given above, and given $\sigma_j^2$, we obtain a fully specified quasi-posterior distribution

$$\prod_{j=1}^{p} \bar{\Pi}_{n,p,j}(\delta, \mathrm{d}\theta, \mathrm{d}\mathsf{q}, \mathrm{d}\rho_{1j}, \mathrm{d}\rho_{2j} | x, \sigma_j^2), \tag{11}$$

where the $j$-th component $\bar{\Pi}_{n,p,j}(\cdot|x, \sigma_j^2)$ can be written as follows. For $\delta \in \Delta_p$, let $\mu_\delta$ be the product measure on $\mathbb{R}^{p-1}$ defined as $\mu_\delta(\mathrm{d}u) = \prod_{j=1}^{p-1} \nu_{\delta_j}(\mathrm{d}u_j)$, where $\nu_0(\mathrm{d}z)$ is the Dirac mass at 0, and $\nu_1(\mathrm{d}z)$ is the Lebesgue measure on $\mathbb{R}$. Then

$$\begin{aligned}\bar{\Pi}_{n,p,j}(\delta, \mathrm{d}\theta, \mathrm{d}\mathsf{q}, \mathrm{d}\rho_{1j}, \mathrm{d}\rho_{2j} | x, \sigma_j^2) \propto q_j(\theta; x) \left(\frac{\mathsf{q}}{1-\mathsf{q}}\right)^{\|\delta\|_1} (1-\mathsf{q})^{d+d^u-1} \\ \times \left(\frac{1}{C_j}\right)^{\|\delta\|_1} e^{-\alpha \frac{\rho_{1j}}{\sigma_j^2}\|\theta\|_1 - (1-\alpha)\frac{\rho_{2j}}{\sigma_j^2}\frac{\|\theta\|_2^2}{2}} \phi(\rho_{1j})\phi(\rho_{2j})\mu_\delta(\mathrm{d}\theta)\mathrm{d}\mathsf{q}\mathrm{d}\rho_{1j}\mathrm{d}\rho_{2j}. \end{aligned} \tag{12}$$

Notice that if, instead of the uniform prior distribution $\phi$, we use a point mass prior distribution for $\rho_{1j}$ and $\rho_{2j}$ in (11), and integrate out $\mathsf{q}$ and $\delta$, we recover exactly (8) where the prior $\Pi_j$ is given by (9) and (10). The quasi-posterior distribution (12) depends on the choice of $\sigma_j^2$. Ideally we would like to set $\sigma_j^2 = 1/\vartheta_{jj}$. However this quantity is unknown. In practice, we suggest choosing $\sigma_j^2$ by empirical Bayes, following Atchadé (2015a) Section 4.3. We explore this approach in the simulations.

---

[3] can be explicitly computed as $C_j = C_\alpha\left(\frac{\rho_{1j}}{\sigma_j^2}, \frac{\rho_{2j}}{\sigma_j^2}\right)$, where, with $\Phi$ denoting the cdf of standard normal distribution, $\mathsf{erfcx}(x) = 2e^{x^2}\Phi(-\sqrt{2}x)$ denoting the scaled complementary error function,

$$C_\alpha(\lambda_1, \lambda_2) \stackrel{\text{def}}{=} \begin{cases} \sqrt{\frac{2\pi}{(1-\alpha)\lambda_2}}\mathsf{erfcx}\left(\frac{\alpha\lambda_1}{\sqrt{2(1-\alpha)\lambda_2}}\right) & \text{if } \alpha \in [0, 1) \\ \frac{2}{\lambda_1} & \text{if } \alpha = 1 \end{cases}.$$

3.2. **Approximate MCMC simulation.** Given $j \in \{1, \ldots, p\}$, sampling from the distribution $\bar{\Pi}_{n,p,j}(\cdot|x)$ given in (12) is a difficult computation task, due to a lack of smoothness in $\theta$, and its trans-dimensional nature[4]. Here we follow the approach developed by the author in Atchadé (2015a), which produces approximate samples from (12) by sampling from its Moreau-Yosida approximation $\bar{\Pi}^{(\gamma)}_{n,p,j}(\delta, \mathrm{d}\theta, \mathrm{d}\mathsf{q}, \mathrm{d}\rho_{1j}, \mathrm{d}\rho_{2j}|x, \sigma_j^2)$. The parameter $\gamma \in (0, 1/4]$ controls the quality of the approximation. It is shown (Atchadé (2015a) Theorem 5) that $\bar{\Pi}^{(\gamma)}_{n,p,j}(\cdot|x, \sigma_j^2)$ converges weakly to $\bar{\Pi}_{n,p,j}(\cdot|x, \sigma_j^2)$ as $\gamma \to 0$. The idea of working with the Moreau-Yosida approximation instead of the distribution itself is attractive because for $\gamma > 0$ fixed, all the probability measures $\bar{\Pi}^{(\gamma)}_{n,p,j}(\delta, \cdot|x, \sigma_j^2)$ for $\delta \in \Delta_p$ are smooth and have densities with respect to the (same) Lebesgue measure $\mathrm{d}\theta\mathrm{d}\mathsf{q}\mathrm{d}\rho_{1j}\mathrm{d}\rho_{2j}$. As a result, one can sample easily from $\bar{\Pi}^{(\gamma)}_{n,p,j}(\cdot|x, \sigma_j^2)$ without any need for trans-dimensional MCMC techniques.

3.3. **Posterior contraction and rate.** Despite the fact that the number of columns ($p$) and the dimension of each column ($p-1$) are both increasing, we will show next that for $n$ reasonably large and for a well-behaved underlying distribution, typical realizations of the quasi-posterior distribution $\check{\Pi}_{n,d}$ given in (8) put most of its probability mass on small neighborhoods of the true value of the parameter.

Given a random sample $X \in \mathbb{R}^{n\times p}$, we shall study the behavior of the random probability measure $\check{\Pi}_{n,p}(\mathrm{d}\theta|X)$ on $\mathbb{R}^{(p-1)\times p}$ as given in (8), for large $n, p$. We focus on the case where the prior distribution is given by (9)-(10), with $\alpha = 1$ (hence $\rho_{2j}$ is irrelevant), and $\rho_{1j}$ fixed. The choice $\alpha = 1$ corresponds to the Laplace prior ($\ell^1$ prior), and is made here mainly for simplicity. We assume below that the rows of $X$ are i.i.d. random variables from a mean-zero Gaussian distribution with precision matrix $\vartheta$.

**H1.** *For some $\vartheta \in \mathcal{M}_p^+$, $X = Z\vartheta^{-1/2}$, where $Z \in \mathbb{R}^{n\times p}$ is a random matrix with i.i.d. standard normal entries.*

From the true precision matrix $\vartheta$, we now form the true value of the parameter $\theta_\star \in \mathbb{R}^{(p-1)\times p}$ towards which $\check{\Pi}_{n,p}$ should converge. For $j = 1, \ldots, p$, $\theta_{\star kj} = -\vartheta_{kj}/\vartheta_{jj}$, for $k = 1, \ldots, j-1$, and $\theta_{\star kj} = -\vartheta_{(k+1)j}/\vartheta_{jj}$, for $k = j, \ldots, p-1$. Let $\delta_\star \in \{0,1\}^{(p-1)\times p}$ be the sparsity structure of $\theta_\star$, defined as $\delta_{\star kj} = \mathbf{1}_{\{|\theta_{\star kj}|>0\}}$. We set

$$s_{\star j} \stackrel{\text{def}}{=} \sum_{k=1}^{p-1} \mathbf{1}_{\{|\theta_{\star kj}|>0\}}, \quad j = 1, \ldots, p \quad \text{and} \quad s_\star \stackrel{\text{def}}{=} \max_{1\le j\le p} s_{\star j}.$$

[4]for two different elements $\delta, \delta'$ of $\Delta_p$, the probability measures $\check{\Pi}_{n,p,j}(\delta, \cdot|x, \sigma_j^2)$ and $\check{\Pi}_{n,p,j}(\delta', \cdot|x, \sigma_j^2)$ are mutually singular

Hence $s_{\star j}$ is the degree of node $j$, and $s_\star$ is the maximum node degree in the undirected graph defined by $\vartheta$.

The asymptotic behavior of $\check{\Pi}_{n,p}$ depends crucially on certain restricted and $m$-sparse eigenvalues of the true precision matrix $\vartheta$, that we introduce next. We set

$$\underline{\kappa} \stackrel{\text{def}}{=} \inf_{\delta\in\{0,1\}^p:\ \|\delta\|_0\le s_\star} \inf\left\{ \frac{u'\vartheta u}{\|u\|_2^2} :\ u\in\mathbb{R}^p,\ u\neq 0,\ \text{s.t.} \sum_{k:\ \delta_k=0} |u_k| \le 7 \sum_{k:\ \delta_k=1} |u_k| \right\}, \tag{13}$$

and for $1\le s\le p$,

$$\underset{\sim}{\kappa}(s) \stackrel{\text{def}}{=} \inf\left\{ \frac{u'\vartheta u}{\|u\|_2^2} :\ u\in\mathbb{R}^p,\ 1\le\|u\|_0\le s \right\},$$

$$\tilde{\kappa}(s) \stackrel{\text{def}}{=} \sup\left\{ \frac{u'\vartheta u}{\|u\|_2^2} :\ u\in\mathbb{R}^p,\ 1\le\|u\|_0\le s \right\}. \tag{14}$$

In the above equations, we convene that $\inf\emptyset = +\infty$, and $\sup\emptyset = 0$. We shall make the following assumption on the prior distribution $\Pi$ on $\mathbb{R}^{(p-1)\times p}$.

**H2.** $\Pi(\mathrm{d}\theta) = \prod_{j=1}^p \Pi_j(\mathrm{d}\theta_{\cdot j})$, *where for each* $j\in\{1,\dots,p\}$ $\Pi_j$ *is of the form described in (9), with* $\alpha_j = 1$, *and*

$$\rho_{1j} = \rho_j \stackrel{\text{def}}{=} \sqrt{\frac{54\tilde{\kappa}(1)}{\vartheta_{jj}} n\log(p)}. \tag{15}$$

*Furthermore, the distribution* $\{\pi_\delta,\ \delta\in\Delta_p\}$ *satisfies:* $\pi_\delta = g_{\|\delta\|_0}\binom{p-1}{\|\delta\|_0}^{-1}$, *for a discrete distribution* $\{g_s,\ 0\le s\le p-1\}$, *for which there exist positive universal constant* $c_1, c_2, c_3, c_4$ *such that*

$$\frac{c_1}{p^{c_3}} g_{s-1} \le g_s \le \frac{c_2}{p^{c_4}} g_{s-1},\ s = 1,\dots,p-1. \tag{16}$$

**Remark 4.** This class of prior distributions was pioneered by Castillo et al. (2014). The example of $\{\pi_\delta\}$ presented in Section 3.1 satisfies (16) with $c_1 = 1/2$, $c_2 = 1$, $c_3 = u$ and $c_4 = u-1$.

Our first result shows that if a minimum sample size requirement is met, and if $\vartheta$ is well-behaved, then typical realizations of $\theta\sim\check{\Pi}_{n,p}(\cdot|X)$ are sparse, with sparsity structure close to the sparsity structure of $\theta_\star$.

**Theorem 5.** *Assume H1 and H2, with* $\underline{\kappa} > 0$. *Suppose also that* $p$ *is large enough so that* $p^{c_4} \ge 8c_2\max(1, 2c_2)$, *with* $c_2, c_4$ *as in H2. For* $1\le j\le p$, *set*

$$\zeta_j = \frac{4}{c_4} + s_{\star j} + \frac{2}{c_4}\left( \frac{\log(4ep)}{\log(p)} + \frac{6912}{\sigma_j^2\vartheta_{jj}} \frac{\tilde{\kappa}(1)}{\underline{\kappa}} + \frac{\sigma_j^2\vartheta_{jj}}{24(\log(p))^2} \frac{\tilde{\kappa}(s_\star)}{\tilde{\kappa}(1)} \right) s_{\star j}.$$

*Then there exist universal finite constants $a_1 > 0$, $a_2 > 0$ such that if*

$$n \geq a_1 s_\star \left(1 + \frac{\tilde{\kappa}(1)}{\underline{\kappa}}\right) \log(p), \tag{17}$$

$$\mathbb{E}\left[\check{\Pi}_{n,p}\left(\left\{\theta \in \mathbb{R}^{(p-1)\times p} : \ \|\theta_{\cdot j}\|_0 \geq \zeta_j \text{ for some } j\right\} \ |X\right)\right] \leq 2\left(\frac{1}{e^{a_2 n}} + \frac{2}{p}\right).$$

*Proof.* See Section 6.1. □

**Remark 6.** In the ideal case where $\sigma_j^2 = 1/\vartheta_{jj}$, and for $n, p$ large, we see that if $\theta \sim \check{\Pi}_{n,p}(\cdot|X)$, then with high probability $\|\theta_{\cdot j}\|_0 < \zeta_j$ for all $j \in \{1, \ldots, p\}$, and

$$\zeta_j \approx \frac{4}{c_4} + s_{\star j} + \frac{2}{c_4}\left(1 + 6912\frac{\tilde{\kappa}(1)}{\underline{\kappa}}\right) s_{\star j}.$$

Hence if the restricted condition number $\tilde{\kappa}(1)/\underline{\kappa}$ remains small, then for large values of $p$, typical realizations of $\check{\Pi}_{n,p}$ are sparse. The large constant 6912 appearing in the theorem is most likely an artifact of the techniques used in the proof, and can probably be improved.

For $\theta \in \mathbb{R}^{(p-1)\times p}$, we set

$$|\!|\!|\theta|\!|\!| \stackrel{\text{def}}{=} \max_{1 \leq j \leq p} \|\theta_{\cdot j}\|_2.$$

**Theorem 7.** *Assume H1, H2, with $\underline{\kappa} > 0$. Let $\zeta_j$ be as in theorem 5. Set $\bar{s} \stackrel{\text{def}}{=} \max_{1\leq j\leq p} \bar{s}_j$, where $\bar{s}_j \stackrel{\text{def}}{=} s_{\star j} + \zeta_j$ if $s_{\star j} > 0$, and $\bar{s}_j \stackrel{\text{def}}{=} 1$ otherwise. Set*

$$\epsilon \stackrel{\text{def}}{=} 12\sqrt{6}\frac{\sqrt{\tilde{\kappa}(1)}}{\underline{\kappa}(\bar{s})}\sqrt{\frac{\bar{s}\log(p)}{n}},$$

*and $M_0 \stackrel{\text{def}}{=} \max\left(96, (4 + c_4(2 + c_3)/2)\max_j \sigma_j^2 \vartheta_{jj}\right)$. Then there exists universal finite constant $a_1 > 0$, $a_2 > 0$ such that if $p \geq \max(24e, 2/c_1)$, $p^{c_4} \geq 8c_2 \max(1, 2c_2)$, and*

$$n \geq a_1 s_\star \left(\frac{\bar{\kappa}(1)}{\underline{\kappa}}\right) \log(p), \quad \text{and} \quad n \geq a_1 \bar{s} \log(p), \tag{18}$$

*then*

$$\mathbb{E}\left[\check{\Pi}_{n,p}\left(\left\{\theta \in \mathbb{R}^{(p-1)\times p} : \ |\!|\!|\theta - \theta_\star|\!|\!| > M_0\epsilon\right\} |X\right)\right] \leq 3\left(\frac{1}{e^{a_2 n}} + \frac{4}{p}\right).$$

*Proof.* See Section 6.2. □

**Remark 8.** Theorem 7 shows that the contraction rate of $\check{\Pi}_{n,p}$ towards $\theta_\star$ in the $|\!|\!|\cdot|\!|\!|$ norm is $O(\epsilon)$. This corresponds to the worst rate among the rates of contraction of

the $p$ linear regression problems performed during the neighborhood selection procedure. This rate is similar to the rate of convergence of the (frequentist) neighborhood selection method of Meinshausen and Buhlmann (2006), which is of order

$$\sqrt{\frac{s_\star \log(p)}{n}}, \tag{19}$$

(see the discussion in Section 3.4 of Ravikumar et al. (2011)). The main difference between (19) and the rate in Theorem 7 is in the dependence on the maximum degree $s_\star$. In the Bayesian case, $s_\star$ is replaced by a worst-case estimate from Theorem 5, namely the largest value that the maximum degree of $\theta \sim \check{\Pi}_{n,p}(\cdot|X)$ can take (with a significant probability).

An interesting difference pointed out in Ravikumar et al. (2011) (Section 3.4), between the neighborhood selection approach and graphical lasso approaches, is that neighborhood selection methods requires a sample size $n$ that scales linearly in $s_\star$, whereas graphical lasso methods require a sample size sample that scales quadratically in $s_\star$. We recover the same dependence on $n$ in Equations (17) and (18) of Theorems 5 and 7, where the sample size scales linearly in $\bar{s}$.

## 4. Numerical experiments

We evaluate the behavior of the quasi-posterior distribution (8) on three simulated datasets. As benchmark, we also report the results obtained using the elastic net estimator

$$\hat{\vartheta}_{\mathsf{glasso}} = \mathsf{Argmin}_{\ \theta\in\mathcal{M}_p^+} \left[ -\log\det\theta + \mathsf{Tr}(\theta S) + \lambda \sum_{i,j} \left( \alpha|\theta_{ij}| + \frac{(1-\alpha)}{2}\theta_{ij}^2 \right) \right],$$

where $S = (1/n)x'x$, $\alpha = 0.9$, and $\lambda > 0$ is a regularization parameter. We choose $\lambda$ by minimizing $-\log\det\left(\hat{\theta}(\lambda)\right) + \mathsf{Tr}(\hat{\theta}(\lambda)S) + \log(n)\sum_{i<j}\mathbf{1}_{\{|\hat{\theta}(\lambda)_{ij}|>0\}}$, over a finite set of values of $\lambda$. We hasten to add that our goal is not to compare the quasi-Bayesian method to graphical `lasso`, since the former utilizes vastly more computing power that the latter. The outputs are also very different, since `Glasso` gives only a point estimate whereas the Bayesian approach produces a full posterior distribution. Rather, we report these numbers as references, to help the reader better understand the behavior of the proposed methodology.

4.1. **Simulation set ups.** We generate a data matrix $x \in \mathbb{R}^{n\times p}$ with i.i.d. rows from $\mathbf{N}(0, \vartheta^{-1})$, $\vartheta \in \mathcal{M}_p^+$. Throughout we set the sample size to $n = 250$, and we consider three settings.

**(a):** $\vartheta$ is generated as in Setting (c) below, but using $p = 100$ nodes.

**(b):** In this case $p = 500$, and we take $\vartheta$ from the R-package `space` based on the work Peng et al. (2009)[5]. These authors have designed a precision matrix $\vartheta$ that is modular with 5 modules of 100 nodes each. Inside each module, there are 3 hubs with degree around 15, and 97 other nodes with degree at most 4. The total number of edges is 587. The resulting partial correlations fall within $(-0.67, -0.10] \cup [0.10, 0.67)$. As explained in Peng et al. (2009), this type of networks are useful models for biological networks.

**(c):** In this case $p = 1,000$, and we build $\vartheta$ as follows. First we generate a symmetric sparse matrix $B$ such that the number of off-diagonal non-zeros entries is roughly $2p$. We magnified the signal by adding 3 to all the non-zeros entries of $B$ (subtracting 3 for negative non-zero entries). Then we set $\vartheta = B + (\epsilon - \lambda_{\min}(B))I_p$, where $\lambda_{\min}(B)$ is the smallest eigenvalue of $B$, with $\epsilon = 1$. In this example, values of the partial correlations are typically in the range $(-0.46, -0.18] \cup [0.18, 0.48)$.

To evaluate the effect of the hyper-parameter $\sigma_j^2$, we report two sets of results. One where $\sigma_j^2 = 1/\vartheta_{jj}$, and another set of results where $\vartheta_{jj}$ is assumed unknown and we select $\sigma_j^2$ from the data, using the cross-validation estimator described in Reid et al. (2013) (see also Atchadé (2015a) Section 4.3).

In order to mitigate the uncertainty in some of the results reported below, we repeat all the MCMC simulations 20 times. Hence, to summarize, for each setting (a), (b), and (c), we generate one precision matrix $\vartheta$. Given $\vartheta$, we generate 20 datasets, and for each dataset, we run two MCMC samplers (one where the $\sigma_j^2$'s are taken as the $1/\vartheta'_{jj}s$, and one where they are estimated from the data).

4.2. **Estimation details.** As explained in Section 3.2, we first approximate the target quasi-posterior

$$\prod_{j=1}^{p} \bar{\Pi}_{n,p,j}(\delta, \mathrm{d}\theta, \mathrm{d}\mathsf{q}, \mathrm{d}\rho_{1j}, \mathrm{d}\rho_{2j} | x, \sigma_j^2)$$

by

$$\prod_{j=1}^{p} \bar{\Pi}^{(\gamma)}_{n,p,j}(\delta, \mathrm{d}\theta, \mathrm{d}\mathsf{q}, \mathrm{d}\phi_{1j}, \mathrm{d}\phi_{2j} | x, \sigma_j^2), \tag{20}$$

where $\bar{\Pi}^{(\gamma)}_{n,p,j}(\cdot|x, \sigma_j^2)$ is the Moreau-Yosida approximation of $\bar{\Pi}_{n,p,j}(\cdot|x, \sigma_j^2)$ given in (12). In all the simulations below, we use $\gamma = 0.2$. We then sample from (20) by parallel computing, each distribution $\bar{\Pi}^{(\gamma)}_{n,p,j}(\cdot|x, \sigma_j^2)$ at the time, and using the MCMC

[5]The precision matrix used here corresponds to the example "Hub network" in Section 3 of Peng et al. (2009). A non-sparse version of $\vartheta$ is attached to the `space` package

sampler developed in Atchadé (2015a). We use a high-performance computer with 100 nodes.

To simulate from $\bar{\Pi}^{(\gamma)}_{n,p,j}(\cdot|x,\sigma_j^2)$ for a given $j$, we run the MCMC sampler for $50,000$ iterations and discard the first $10,000$ iterations as burn-in. We use Geweke's diagnostic test on the remaining samples to test for convergence using the negative pseudo-log-likelihood function $\theta \mapsto \frac{1}{2\sigma_j^2}\|x_{\cdot j} - x^{(j)}\theta\|_2^2$. All the samplers passed the test. This suggests that $50,000$ is a reasonably large number of iterations for these examples.

From the MCMC output, we estimate the structure $\delta \in \{0,1\}^{p\times p}$ as follows. We set the diagonal of $\delta$ to one, and for each off-diagonal entry $(i,j)$ of $\delta$, we estimate $\delta_{ij}$ as equal to 1 if the sample average estimate of $\delta_{ij}$ (from the $j$-th chain) and the sample average estimate of $\delta_{ji}$ (from the $i$-th chain) are both larger than 0.5. Otherwise $\delta_{ij} = 0$. Obviously, other symmetrization rules could be adopted.

Given the estimate $\hat{\delta}$ say, of $\delta$, we estimate $\vartheta \in \mathbb{R}^{p\times p}$ as follows. We set the diagonal of $\vartheta$ to $(1/\sigma_j^2)$. For $i \neq j$, if $\hat{\delta}_{ij} = 0$, we set $\vartheta_{ij} = \vartheta_{ji} = 0$. Otherwise we estimate $\vartheta_{ij} = \vartheta_{ji}$ as $0.5(-1/\sigma_j^2)\bar{\vartheta}_{ij} + 0.5(-1/\sigma_i^2)\bar{\vartheta}_{ji}$, where $\bar{\vartheta}_{ij}$ (resp. $\bar{\vartheta}_{ji}$) is the Monte Carlo sample average estimate of $\vartheta_{ij}$ from the $j$-th chain (resp. $i$-th chain). For all the off-diagonal components $(i,j)$ such that $\hat{\delta}_{ij} = 1$, we also produce a 95% posterior interval by taking the union of the 95% posterior intervals from the $i$-th and $j$-th chains. When $\hat{\delta}_{ij} = 0$, we set the confidence interval to $\{0\}$.

4.3. **Results.** We look at the performance of the method by computing the relative Frobenius norm, the sensitivity and the precision of the estimated matrix (as obtained above). These quantities are defined respectively as

$$\mathcal{E} = \frac{\|\hat{\vartheta} - \vartheta\|_{\mathsf{F}}}{\|\vartheta\|_{\mathsf{F}}}, \quad \mathsf{SEN} = \frac{\sum_{i<j} \mathbf{1}_{\{|\vartheta_{ij}|>0\}}\mathbf{1}_{\{\mathsf{sign}(\hat{\vartheta}_{ij})=\mathsf{sign}(\vartheta_{ij})\}}}{\sum_{i<j}\mathbf{1}_{\{|\vartheta_{ij}|>0\}}};$$
$$\text{and} \quad \mathsf{PREC} = \frac{\sum_{i<j} \mathbf{1}_{\{|\hat{\vartheta}_{ij}|>0\}}\mathbf{1}_{\{\mathsf{sign}(\hat{\vartheta}_{ij})=\mathsf{sign}(\vartheta_{ij})\}}}{\sum_{i<j}\mathbf{1}_{\{|\hat{\vartheta}_{ij}|>0\}}}. \quad (21)$$

We average these statistics over the 20 simulations replications. We compute also the same quantities for the elastic net $\hat{\vartheta}_{\mathsf{glasso}}$. These results are reported in Table 1-3. These results suggest that the quasi-Bayesian procedure has good contraction properties in the Frobenius norm (and hence in the $L_{\infty,2}$ norm). The results also suggest that the quasi-Bayesian procedure is not very sensitive (it has a high false negative rate), but has excellent precision (it has a very low false positive rate), even with $p = 1,000$. The same conclusion seems to hold across all three network settings considered in the simulations.

Another interesting point to notice from these results is that there seems to be little difference between the results where $\vartheta_{jj}$ is assumed known and the results where $\vartheta_{jj}$ is estimated from the data.

In a typical use of the method in the applications, one would run the MCMC sampler only once, and compute the posterior estimate, and confidence intervals, for instance as in Section 4.2. We show one such output. In Setting (a), where $p = 100$, we plot on Figure 1 all the 95% confidence intervals for all the off-diagonal elements $\vartheta_{ij}$ of $\vartheta$, obtained from one MCMC run. We also add a dot to the confidence interval line to represent the true value of $\vartheta_{ij}$. The results seem consistent with the results in Table 1.

| | $\vartheta_{jj}^2$ known | Empirical Bayes | `Glasso` |
|---|---|---|---|
| Relative Error ($\mathcal{E}$ in %) | 19.2 | 21.6 | 63.1 |
| Sensitivity (SEN in %) | 68.4 | 69.0 | 40.5 |
| Precision (PREC in %) | 100.0 | 100.0 | 74.9 |

TABLE 1. Table showing the relative error, sensitivity and precision (as defined in (21)) for Setting (a), with $p = 100$ nodes. Based on 20 simulation replications. Each MCMC run is $5 \times 10^4$ iterations.

| | $\vartheta_{jj}^2$ known | Empirical Bayes | `Glasso` |
|---|---|---|---|
| Relative Error ($\mathcal{E}$ in %) | 23.1 | 26.2 | 45.2 |
| Sensitivity (SEN in %) | 44.6 | 45.4 | 87.9 |
| Precision (PREC in %) | 100 | 99.9 | 56.1 |

TABLE 2. Table showing the relative error, sensitivity and precision (as defined in (21)) for Setting (b), with $p = 500$ nodes. Based on 20 simulation replications. Each MCMC run is $5 \times 10^4$ iterations.

| | $\vartheta_{jj}^2$ known | Empirical Bayes | `Glasso` |
|---|---|---|---|
| Relative Error ($\mathcal{E}$ in %) | 30.8 | 35.2 | 66.9 |
| Sensitivity (SEN in %) | 16.3 | 16.4 | 6.6 |
| Precision (PREC in %) | 99.9 | 99.8 | 94.7 |

TABLE 3. Table showing the relative error, sensitivity and precision (as defined in (21)) for Setting (c), with $p = 1,000$ nodes. Based on 20 simulation replications. Each MCMC run is $5 \times 10^4$ iterations.

## 5. Some concluding remarks

We have developed in this work a quasi-Bayesian methodology for inferring high-dimensional Gaussian graphical models by neighborhood selection. We have shown by examples that, using a high-performance computer systems with multiple cores, the method can fit Gaussian graphical models at a scale unmatched by existing Bayesian methodologies. The general discussion in Section 2 also shows that the method can be easily extended to handle other classes of graphical models. We have studied the asymptotic behavior of the method in the Gaussian case, and showed that for sparse and well-behaved problems, the quasi-posterior distribution concentrates around the true value even in setting where $p$ exceeds $n$. One important direction for future work is the extension of the methodology to estimate the scale parameters $\vartheta_{jj}$ jointly, as part of the Bayesian modeling. This extension raises several difficulties, in terms of the computations (the approximation scheme of Atchadé (2015a) cannot readily handle such cases), but also in terms of the Bayesian asymptotics.

## 6. Proof of Theorem 5 and Theorem 7

Similar results have been derived recently for the linear regression model by Castillo et al. (2014), and by the author in Atchadé (2015b). Therefore, a natural strategy to proof Theorem 5 and Theorem 7 is to reduce the problem to a corresponding problem in a linear regression model. In the details, we will rely on the behavior of some restricted and $m$-sparse eigenvalues concepts that we introduce first. For $z \in \mathbb{R}^{n\times q}$, for some $q \geq 1$, and for $s \geq 1$, we define

$$\underline{\kappa}(s,z) \stackrel{\text{def}}{=} \inf_{\delta\in\{0,1\}^q:\ \|\delta\|_0\leq s} \inf\left\{ \frac{\theta'(z'z)\theta}{n\|\theta\|_2^2} :\ \theta\in\mathbb{R}^q,\ \theta\neq 0,\ \sum_{k:\ \delta_k=0}|\theta_k| \leq 7\sum_{k:\ \delta_k=1}|\theta| \right\},$$

and

$$\underset{\sim}{\kappa}(s,z) \stackrel{\text{def}}{=} \inf\left\{ \frac{\theta'(z'z)\theta}{n\|\theta\|_2^2} :\ \theta\in\mathbb{R}^q,\ 1\leq\|\theta\|_0\leq s \right\},$$

$$\tilde{\kappa}(s,z) \stackrel{\text{def}}{=} \sup\left\{ \frac{\theta'(z'z)\theta}{n\|\theta\|_2^2} :\ \theta\in\mathbb{R}^q,\ 1\leq\|\theta\|_0\leq s \right\}.$$

In the above definition, we convene that $\inf\emptyset = +\infty$, and $\sup\emptyset = 0$. We are interested in the behavior of $\underline{\kappa}(s_\star, X)$, $\underset{\sim}{\kappa}(s, X)$ and $\tilde{\kappa}(s, X)$, when $X$ is the random matrix obtained from assumption H1. We will use the following result taken from Raskutti et al. (2010) Theorem 1, and Rudelson and Zhou (2013) Theorem 3.2, which relates the behavior of $\underline{\kappa}(s_\star, X)$, $\underset{\sim}{\kappa}(s, X)$ and $\tilde{\kappa}(s, X)$ to the corresponding term $\underline{\kappa}$, $\underset{\sim}{\kappa}(s)$ and $\tilde{\kappa}(s)$ of the true precision matrix $\vartheta$ introduced in (13)-(14).

**Lemma 9.** *Assume H1. Then there exists finite universal constant $a_1 > 0$, $a_2 > 0$ such that for the following hold.*

*(1) If $\underline{\kappa} > 0$, then for all $n \geq a_1 \frac{\tilde{\kappa}(1)}{\underline{\kappa}} s_\star \log(p)$*

$$\mathbb{P}\left[64\underline{\kappa}(s_\star, X) < \underline{\kappa}\right] \leq e^{-a_2 n}.$$

*(2) Let $1 \leq s \leq p$ be such that $\underset{\sim}{\kappa}(s) > 0$, then for all $n \geq a_1 s \log(p)$,*

$$\mathbb{P}\left[4\underset{\sim}{\kappa}(s, X) < \underset{\sim}{\kappa}(s) \;\; or \;\; 4\tilde{\kappa}(s, X) > 9\tilde{\kappa}(s)\right] \leq e^{-a_2 n}.$$

6.1. **Proof of Theorem 5.** We have

$$\check{\Pi}_{n,p}(\mathrm{d}\theta|X) = \prod_{j=1}^{p} \check{\Pi}_{n,p,j}(\mathrm{d}\theta_{\cdot j}|X),$$

where for $j \in \{1, \ldots, p\}$, and using that $\alpha = 1$ in (10), $\check{\Pi}_{n,p,j}(\mathrm{d}\theta_{\cdot j}|X)$ is given by

$$\check{\Pi}_{n,p,j}(\mathrm{d}u|X) \propto q_j(u; X) \sum_{\delta \in \Delta_p} \pi_\delta \left(\frac{\rho_j}{2\sigma_j^2}\right)^{\|\delta\|_1} e^{-\frac{\rho_j}{\sigma_j^2}\|u\|_1} \mu_\delta(\mathrm{d}u), \tag{22}$$

and

$$\log q_j(u; X) = -\frac{1}{2\sigma_j^2}\|X_{\cdot j} - X^{(j)}u\|_2^2.$$

For $q \geq 1$, we define

$$\mathcal{G}_{n,q} \stackrel{\text{def}}{=} \left\{ z \in \mathbb{R}^{n \times q} : \; \tilde{\kappa}(s_\star, z) \leq \frac{9}{4}\tilde{\kappa}(s_\star), \;\; \tilde{\kappa}(1, z) \leq \frac{9}{4}\tilde{\kappa}(1), \;\; \text{and} \;\; \underline{\kappa}(s_\star, z) \geq \frac{\underline{\kappa}}{64} \right\}.$$

For any $k_j \geq 0$, we start by noting that

$$\mathbb{E}\left[\check{\Pi}_{n,p}\left(\left\{\theta \in \mathbb{R}^{(p-1)\times p} : \; \|\theta_{\cdot j}\|_0 \geq k_j, \;\; \text{for some } j\right\}|X\right)\right] \leq \mathbb{P}(X \notin \mathcal{G}_{n,p}) + \sum_{j=1}^{p} \mathbb{E}\left[\mathbf{1}_{\mathcal{G}_{n,p}}(X)\check{\Pi}_{n,p,j}\left(\mathcal{A}_j|X\right)\right].$$

where $\mathcal{A}_j \stackrel{\text{def}}{=} \{u \in \mathbb{R}^{p-1} : \; \|u\|_0 \geq k_j\}$. We notice that if $X \in \mathcal{G}_{n,p}$, then $X^{(j)} \in \mathcal{G}_{n,p-1}$ for any $1 \leq j \leq p$. We recall that the notation $X^{(j)}$ denotes the matrix obtained by removing the $j$ column of $X$. Hence

$$\mathbb{E}\left[\mathbf{1}_{\mathcal{G}_{n,p}}(X)\check{\Pi}_{n,p,j}\left(\mathcal{A}_j|X\right)\right] \leq \mathbb{E}\left[\mathbf{1}_{\mathcal{G}_{n,p-1}}(X^{(j)})\check{\Pi}_{n,p,j}\left(\mathcal{A}_j|X\right)\right] = \mathbb{E}\left[\mathbf{1}_{\mathcal{G}_{n,p-1}}(X^{(j)})\mathbb{E}\left(\check{\Pi}_{n,p,j}\left(\mathcal{A}_j|X\right)|X^{(j)}\right)\right].$$

We conclude that

$$\mathbb{E}\left[\check{\Pi}_{n,p}\left(\left\{\theta\in\mathbb{R}^{(p-1)\times p}:\ \|\theta_{\cdot j}\|_0\geq k_j,\ \text{ for some } j\right\}|X\right)\right] \leq \mathbb{P}(X\notin\mathcal{G}_{n,p})+\sum_{j=1}^{p}\mathbb{E}\left[\mathbf{1}_{\mathcal{G}_{n,p-1}}(X^{(j)})T_j\right], \quad (23)$$

where

$$T_j=\mathbb{E}\left(\check{\Pi}_{n,p,j}\left(\mathcal{A}_j|X\right)|X^{(j)}\right).$$

The main idea of the proof is to notice that $T_j$ is an expected quasi-posterior probability in the linear regression model $X_{\cdot j}=X^{(j)}\beta+\eta$, where $\eta\sim\mathbf{N}(0,(1/\vartheta_{jj})I_n)$. Therefore, by a similar argument and similar calculations as in the proof of Theorem 13 of Atchadé (2015b), we have

$$T_j\leq 2p\exp\left(-\frac{\vartheta_{jj}\rho_j^2}{8\max_{k\neq j}\|X_{\cdot k}\|_2^2}\right) + 2(4^{s_{\star j}})\left(1+\frac{\sigma_j^2L_j}{\rho_j^2}\right)^{s_{\star j}}e^{\frac{2\rho_j^2s_{\star j}}{\tau_j\sigma_j^2}}\binom{p-1}{s_{\star j}}\left(\frac{4c_2}{p^{c_4}}\right)^{k_j-s_{\star j}}, \quad (24)$$

where $L_j=n\tilde{\kappa}(s_\star,X^{(j)})$, and $\tau_j=n\underset{\sim}{\kappa}(s_\star,X^{(j)})$. Since $\max_{k\neq j}\|X_{\cdot k}\|_2^2=n\tilde{\kappa}(1,X^{(j)})$, for $X^{(j)}\in\mathcal{G}_{n,p-1}$, it is easy to see that the first term on the right-hand side of (24) is bounded by

$$2p\exp\left(-\frac{\vartheta_{jj}\rho_j^2}{18n\tilde{\kappa}(1)}\right)=\frac{2}{p^2},$$

where the equality follows from the choice of $\rho_j$. Using the fact that for $X^{(j)}\in\mathcal{G}_{n,p-1}$, we have $L_j\leq(9/4)n\tilde{\kappa}(s_\star)$, $\tau_j\geq(1/64)n\underset{\sim}{\kappa}$, it is easy to show that the second term on the right-hand side of (24) is bounded by

$$2\exp\left[s_{\star j}\log(p)\left(\frac{6912}{\sigma_j^2\vartheta_{jj}}\frac{\tilde{\kappa}(s_\star)}{\underset{\sim}{\kappa}}+\frac{\sigma_j^2\vartheta_{jj}}{24(\log(p))^2}\frac{\tilde{\kappa}(s_{\star j})}{\tilde{\kappa}(1)}+\frac{\log(4ep)}{\log(p)}\right)-\frac{c_4}{2}(k_j-s_{\star j})\log(p)\right].$$

With $k_j=\zeta_j$ as given in the statement of the theorem, this latter expression is bounded by $2/(p^2)$. This conclude the proof.

6.2. **Proof of Theorem 7.** We use the same approach as above. We define $\bar{s}_j=s_{\star j}+\zeta_j$ ($\bar{s}_j=1$ if $s_{\star j}=0$), and $\bar{s}=\max_j\bar{s}_j$, and we set

$$\mathcal{G}_{n,q}\overset{\text{def}}{=}\left\{z\in\mathbb{R}^{n\times q}:\ \tilde{\kappa}(s_\star,z)\leq\frac{9}{4}\tilde{\kappa}(s_\star),\ \text{ and }\ \underset{\sim}{\kappa}(\bar{s},z)\geq\frac{1}{4}\underset{\sim}{\kappa}(\bar{s})\right\}.$$

We also define $\mathcal{U} \stackrel{\text{def}}{=} \{\theta \in \mathbb{R}^{(p-1)\times p} : \ \|\theta_{\cdot j} - \theta_{\star\cdot j}\|_2 > \epsilon_j, \ \text{ for some } j\}$, $\bar{\mathcal{U}} \stackrel{\text{def}}{=} \mathcal{U} \cap \{\theta \in \mathbb{R}^{(p-1)\times p} : \ \|\theta_{\cdot j} - \theta_{\star\cdot j}\|_0 \leq s_{\star j} + \zeta_j \ \text{ for all } j\}$, and

$$\check{\Pi}_{n,p}(\mathcal{U}|X) \leq \check{\Pi}_{n,p}\left(\{\theta \in \mathbb{R}^{(p-1)\times p} : \ \|\theta_{\cdot j} - \theta_{\star\cdot j}\|_0 > s_{\star j} + \zeta_j \ \text{ for some } j\}|X\right) + \mathbf{1}_{\mathcal{G}_{n,p}^c}(X) + \mathbf{1}_{\mathcal{G}_{n,p}}(X)\check{\Pi}_{n,p}\left(\bar{\mathcal{U}}|X\right). \quad (25)$$

If for some $j$, $\|\theta_{\cdot j} - \theta_{\star\cdot j}\|_0 > s_{\star j} + \zeta_j$, then we necessarily have $\|\theta_{\cdot j}\|_0 > \zeta_j$. Therefore, by Theorem 5, we have:

$$\mathbb{E}\left[\check{\Pi}_{n,p}\left(\{\theta \in \mathbb{R}^{(p-1)\times p} : \ \|\theta_{\cdot j} - \theta_{\star\cdot j}\|_0 > s_{\star j} + \zeta_j \ \text{ for some } j\}|X\right)\right] \leq \frac{2}{e^{a_2 n}} + \frac{4}{p}. \quad (26)$$

By Lemma 9, for $n \geq a_1 \bar{s}\log(p)$,

$$\mathbb{E}\left[\mathbf{1}_{\mathcal{G}_{n,p}^c}(X)\right] = \mathbb{P}\left[X \notin \mathcal{G}_{n,p}\right] \leq \frac{1}{e^{a_2 n}}. \quad (27)$$

It remains to control the last term on the right-hand side of (25). To do so, we note that if $X \in \mathcal{G}_{n,p}$, then $X^{(j)} \in \mathcal{G}_{n,p-1}$ for all $1 \leq j \leq p$. Hence

$$\begin{aligned}\mathbb{E}\left[\mathbf{1}_{\mathcal{G}_{n,p}}(X)\check{\Pi}_{n,p}\left(\bar{\mathcal{U}}|X\right)\right] &\leq \sum_{j=1}^{p}\mathbb{E}\left[\mathbf{1}_{\mathcal{G}_{n,p-1}}(X^{(j)})\check{\Pi}_{n,p,j}(\mathcal{A}_j|X)\right] \\ &\leq \sum_{j=1}^{p}\mathbb{E}\left[\mathbf{1}_{\mathcal{G}_{n,p-1}}(X^{(j)})\mathbb{E}\left(\check{\Pi}_{n,p,j}(\mathcal{A}_j|X)|X^{(j)}\right)\right], \quad (28)\end{aligned}$$

where $\mathcal{A}_j \stackrel{\text{def}}{=} \{u \in \mathbb{R}^{p-1} : \ \|u - \theta_{\star\cdot j}\|_2 > \epsilon_j, \ \text{ and } \ \|u - \theta_{\star\cdot j}\|_0 \leq \bar{s}_j\}$. As in the proof of Theorem 5, we note that under the conditional distribution of $X_{\cdot j}$ given $X^{(j)}$, the term $\check{\Pi}_{n,p,j}(\mathcal{A}_j|X)$ can be viewed as the posterior distribution in the linear regression model $X_{\cdot j} = X^{(j)}\beta + \eta$, where $\eta \sim \mathbf{N}(0, (1/\vartheta_{jj})I_n)$. Therefore, by proceeding as in the proof of Theorem 13 of Atchadé (2015b), and for any constant $M_0 \geq 96$, we have

$$\begin{aligned}\mathbb{E}\left(\check{\Pi}_{n,p,j}(\mathcal{A}_j|X)|X^{(j)}\right) &\leq 2p\exp\left(-\frac{\vartheta_{jj}\rho_j^2}{8\max_{k\neq j}\|X_{\cdot k}\|_2^2}\right) \\ &+ e^{\bar{s}_j\log(24pe)}\frac{e^{-\frac{M_0^2\tau_j\bar{\epsilon}_j^2}{32}}}{1-e^{-\frac{M_0^2\tau_j\bar{\epsilon}_j^2}{32}}} + 2\binom{p}{s_{\star j}}\left(\frac{p^{c_3}}{c_1}\right)^{s_{\star j}}\left(1+\frac{\sigma_j^2 L_j}{\rho_j^2}\right)^{s_{\star j}}\frac{e^{-\frac{M_0^2\tau_j\bar{\epsilon}_j^2}{64}}}{1-e^{-\frac{M_0^2\tau_j\bar{\epsilon}_j^2}{64}}}, \quad (29)\end{aligned}$$

where $\bar{\epsilon}_j = \frac{\rho_j \bar{s}_j^{1/2}}{\tau_j}$, $\tau_j = n\underset{\sim}{\kappa}(\bar{s}_j, X^{(j)})$, and $L_j = n\tilde{\kappa}(s_{\star j}, X^{(j)})$. As seen in the proof of Theorem 5, the first term on the right-hand side of (29) is upper bounded by $2/p^2$.

We have

$$\frac{M_0^2\tau_j\bar{\epsilon}_j^2}{32} \geq \left(\frac{54M_0^2}{32}\frac{1}{\sigma_j^2\vartheta_{jj}}\right)\bar{s}_j\log(p).$$

Hence for $p \geq 24e$, and $\frac{54M_0^2}{32}\frac{1}{\sigma_j^2\vartheta_{jj}} \geq 4$, the second term on the right-hand side of (29) is also upper bounded by $2/p^2$. For $\frac{54M_0^2}{32}\frac{1}{\sigma_j^2\vartheta_{jj}} \geq 4$, the third term is upper bounded by

$$4\exp\left[s_{\star j}\log(p)\left(2+c_3+\frac{\sigma_j^2\vartheta_{jj}}{24(\log(p)^2)}\frac{\tilde{\kappa}(s_{\star j})}{\tilde{\kappa}(1)}\right)-\frac{54M_0^2}{64}\frac{1}{\sigma_j^2\vartheta_{jj}}\bar{s}_j\log(p)\right] \leq \frac{2}{p^2},$$

by choosing $\frac{54M_0^2}{64}\frac{1}{\sigma_j^2\vartheta_{jj}} \geq 2+\frac{c_4}{2}(2+c_3)$. This conclude the proof.

## Acknowledgements

The author would like to thank Shuheng Zhou for very helpful conversations. This work is partially supported by the NSF, grants DMS 1228164 and DMS 1513040.

## References

Atay-Kayis, A. and Massam, H. (2005). A Monte Carlo method for computing the marginal likelihood in nondecomposable Gaussian graphical models. *Biometrika* **92** 317–335.

Atchadé, Y. F. (2015a). A Moreau-Yosida approximation scheme for high-dimensional quasi-posterior distributions. *ArXiv e-prints* .

Atchadé, Y. F. (2015b). On the contraction properties of some high-dimensional quasi-posterior distributions. *ArXiv e-prints* .

Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *J. Roy. Statist. Soc. Ser. B* **36** 192–236.

Bühlmann, P. and van de Geer, S. (2011). *Statistics for high-dimensional data*. Springer Series in Statistics, Springer, Heidelberg. Methods, theory and applications.

Carvalho, C. M., Polson, N. G. and Scott, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika* **97** 465–480.

Castillo, I., Schmidt-Hieber, J. and van der Vaart, A. W. (2014). Bayesian linear regression with sparse priors. *ArXiv e-prints* .

Cheng, J., Levina, E. and Zhu, J. (2013). High-dimensional Mixed Graphical Models. *ArXiv e-prints* .

Chernozhukov, V. and Hong, H. (2003). An MCMC approach to classical estimation. *J. Econometrics* **115** 293–346.

Dobra, A., Lenkoski, A. and Rodriguez, A. (2011). Bayesian inference for general Gaussian graphical models with application to multivariate lattice data. *J. Amer. Statist. Assoc.* **106** 1418–1433.

Fearnhead, P. and Prangle, D. (2010). Constructing summary statistics for approximate bayesian computation: Semi-automatic abc. *Technical Report, Lancaster University, UK* .

Hastie, T., Tibshirani, R. and Wainwright, M. (2015). *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman and Hall/CRC.

Higham, N. J. (1988). Computing a nearest symmetric positive semidefinite matrix. *Linear Algebra and its Applications* **103** 103 – 118.

Kato, K. (2013). Quasi-Bayesian analysis of nonparametric instrumental variables models. *Ann. Statist.* **41** 2359–2390.

Khondker, Z. S., Zhu, H., Chu, H., Lin, W. and Ibrahim, J. G. (2013). The Bayesian covariance lasso. *Stat. Interface* **6** 243–259.

Lenkoski, A. and Dobra, A. (2011). Computational aspects related to inference in Gaussian graphical models with the G-Wishart prior. *J. Comput. Graph. Statist.* **20** 140–157. Supplementary material available online.

Li, C. and Jiang, W. (2014). Model Selection for Likelihood-free Bayesian Methods Based on Moment Conditions: Theory and Numerical Examples. *ArXiv e-prints* .

Lyne, A.-M., Girolami, M., Atchade, Y., Strathmann, H. and Simpson, D. (2013). On Russian Roulette Estimates for Bayesian inference with doubly-intractable Likelihoods. *ArXiv e-prints* .

Meinshausen, N. and Buhlmann, P. (2006). High-dimensional graphs with the lasso. *Annals of Stat.* **34** 1436–1462.

Mukherjee, S. and Speed, T. P. (2008). Network inference using informative priors. *Proceedings of the National Academy of Sciences* **105** 14313–14318.

Murray, I., Ghahramani, Z. and MacKay, D. (2006). MCMC for doubly-intractable distributions. *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence UAI06* 359–366.

Park, T. and Casella, G. (2008). The Bayesian lasso. *J. Amer. Statist. Assoc.* **103** 681–686.

Peng, J., Wang, P., Zhou, N. and Zhu, J. (2009). Partial correlation estimation by joint sparse regression models. *Journal of the American Statistical Association* **104** 735–746.

Peterson, C., Stingo, F. C. and Vannucci, M. (2015). Bayesian inference of multiple gaussian graphical models. *Journal of the American Statistical Association* **110** 159–174.

Raskutti, G., Wainwright, M. J. and Yu, B. (2010). Restricted eigenvalue properties for correlated gaussian designs. *J. Mach. Learn. Res.* **11** 2241–2259.

Ravikumar, P., Wainwright, M. J., Raskutti, G. and Yu, B. (2011). High-dimensional covariance estimation by minimizing $\ell_1$-penalized log-determinant divergence. *Electron. J. Stat.* **5** 935–980.

Reid, S., Tibshirani, R. and Friedman, J. (2013). A Study of Error Variance Estimation in Lasso Regression. *ArXiv e-prints* .

Rudelson, M. and Zhou, S. (2013). Reconstruction from anisotropic random measurements. *IEEE Trans. Inf. Theor.* **59** 3434–3447.

Salakhutdinov, R. and Hinton, G. E. (2009). Deep boltzmann machines. *Journal of Machine Learning Research - Proceedings Track* **5** 448–455.

Sun, T. and Zhang, C.-H. (2013). Sparse Matrix Inversion with Scaled Lasso. *Journal of Machine Learning Research* **14** 3385–3418.

Wang, H. (2012). Bayesian graphical lasso models and efficient posterior computation. *Bayesian Anal.* **7** 867–886.

Wang, H. and Li, S. Z. (2012). Efficient gaussian graphical model determination under g-wishart prior distributions. *Electron. J. Statist.* **6** 168–198.

Yang, E., Ravikumar, P., Allen, G. I., Baker, Y., Wan, Y.-W. and Liu, Z. (2014). A General Framework for Mixed Graphical Models. *ArXiv e-prints* .

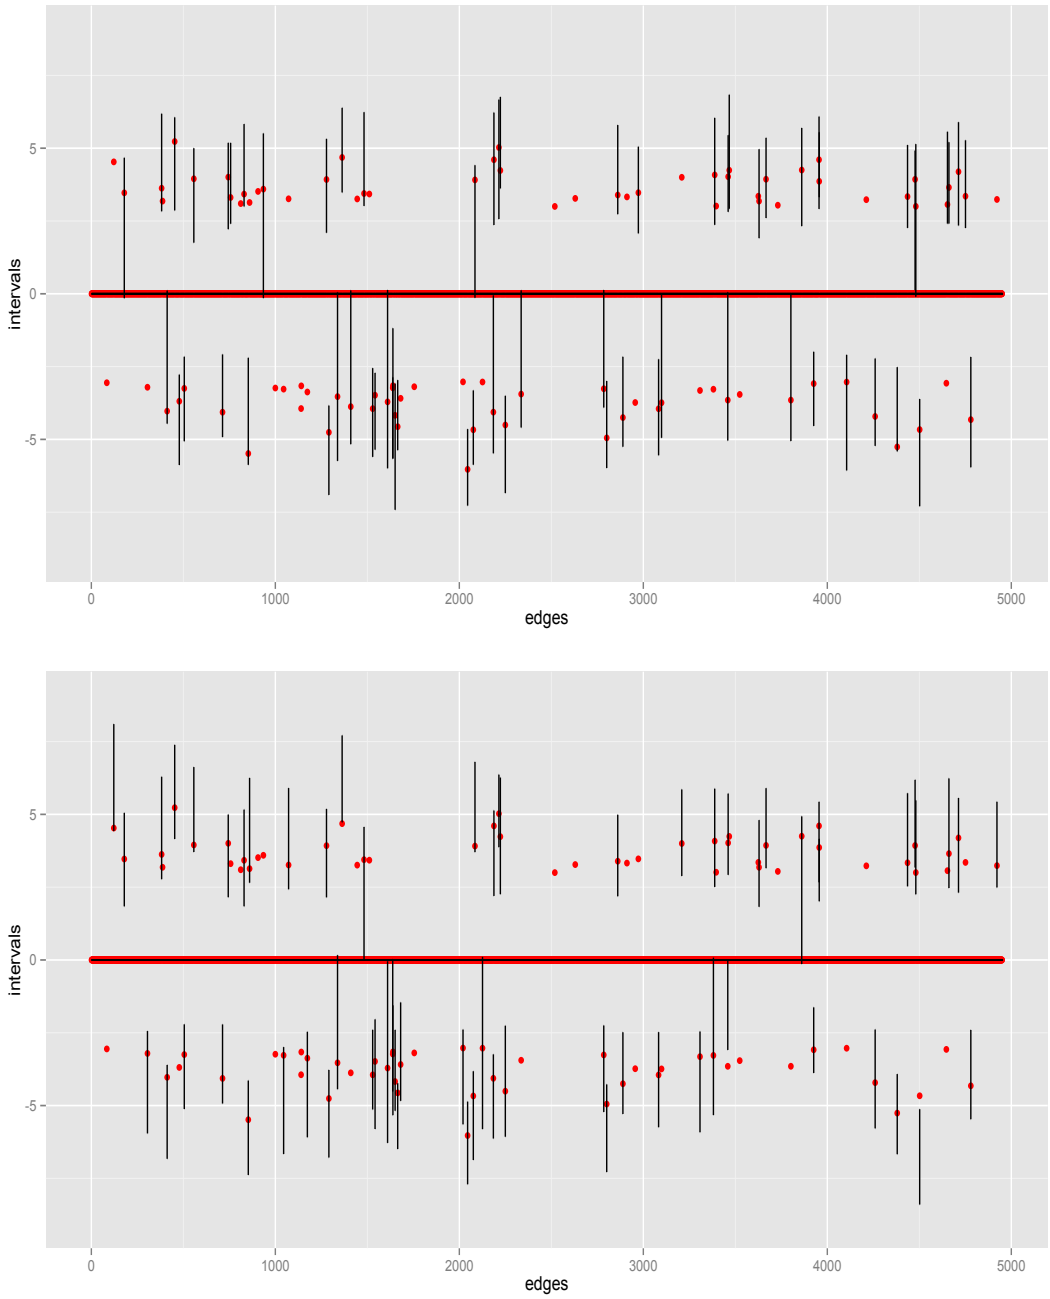


FIGURE 1. Figure showing the confidence interval bars (obtained from one MCMC run), for the non-diagonal entries of $\vartheta$ in Setting (a). The dots represent the true values. Top graph is the case where the $\vartheta_{jj}$'s are assumed known.